

Multiplicative weights, equalizers, and $\mathbf{P}=\mathbf{PPAD}$

Ioannis Avramopoulos

We show that, by using multiplicative weights in a game-theoretic thought experiment (and an important convexity result on the composition of multiplicative weights with the relative entropy function), a symmetric bimatrix game (that is, a bimatrix matrix wherein the payoff matrix of each player is the transpose of the payoff matrix of the other) either has an interior symmetric equilibrium or there is a pure strategy that is weakly dominated by some mixed strategy. Weakly dominated pure strategies can be detected and eliminated in polynomial time by solving a linear program. Furthermore, interior symmetric equilibria are a special case of a more general notion, namely, that of an “equalizer,” which can also be computed efficiently in polynomial time by solving a linear program. An elegant “symmetrization method” of bimatrix games [Jurg et al., 1992] and the well-known **PPAD**-completeness results on equilibrium computation in bimatrix games [Daskalakis et al., 2009, Chen et al., 2009] imply then the compelling $\mathbf{P} = \mathbf{PPAD}$.

1 Introduction

THIS paper is an intellectual exercise using multiplicative weights. There are many variants and our results are perhaps not strongly bound to any particular multiplicative weights algorithm, but our favorite is Hedge [Freund and Schapire, 1997, 1999], which we use as a friendly companion to indulge in our intellectual pursuit. Our motivation is to compute Nash equilibria in polynomial time, noting up front that we do not necessarily use Hedge to compute them—for example, we can show as a simple implication of stronger properties that Hedge diverges from the equilibrium of rock-paper-scissors—but we use Hedge to prove some very nice equilibrium properties as well some properties on the structure of strategic games. The fact that Hedge does not necessarily converge is not necessarily bad—we believe it can be used as a powerful model of “evolutionary behavior” (in the sense of dynamical systems as well as that of biology) to explore “terrains” whether physical or conceptual. This paper actually amounts to some notion of proof of such a claim.

1.1 OUR MAIN RESULT

Our main result is to refute much of the folk wisdom on algorithmic game theory (and perhaps complexity theory in general for that matter) showing $\mathbf{P} = \mathbf{PPAD}$.

1.2 OUR DETAILED QUESTIONS

There are many ways to understand the mathematics of this paper, but one of the most fruitful is as a thought experiment: *Can multiplicative weights serve as game theory’s foundation?* The dominant approach to understanding game theory is by means of the *best response*, the *Nash equilibrium*, and how the two are related—the Nash equilibrium is a profile against which no player has an incentive to deviate or, equivalently, it is a fixed point of the game’s *best response correspondence*. The major problem with this approach is that alternating best responses may oscillate (or worse) and, therefore, game theory’s “natural dynamic” does not correspond to an equilibrium computation algorithm. Nevertheless, the *concept* of the best response dynamic is very helpful as it can be used

to prove that a Nash equilibrium always exists which is soothing in founding a theory of strategic behavior based on the Nash equilibrium. So let us go back to our own question:

1.3 OUR DETAILED RESULTS AND TECHNIQUES

Best response dynamics are very problematic because they are *discontinuous* (and, therefore, often oscillatory or worse). One very prolific way of viewing multiplicative weights is as a *continuous approximation* of best responses, a point which we support by a lemma: Multiplicative weights are a continuous *better response dynamic* (for arbitrary values of the *learning rate* parameter). We do not use this lemma in any of our results and we give it only as an illustration of this important idea. Continuity and the better response property are important in a variety of ways not all of which we explore here. However, our main result crucially relies on this continuity property along with the continuity of the *relative entropy function* (known also as *Kullback-Leibler divergence*) and the elementary result in analysis that the composition of continuous functions is continuous.

1.3.1 Our main overarching strategy

In contrast to the majority of contributions on equilibrium computation focusing on bimatrix games, we focus instead on the more particular class of *symmetric bimatrix games* and even more specifically with the computation of *symmetric equilibria* in this class of games (that is, equilibria wherein both players use the same strategy). Symmetric equilibria in symmetric N -person (and, therefore, also 2-person) games are guaranteed to exist [Nash, 1951]. Symmetric equilibria are also a foundational concept in *evolutionary game theory* [Maynard Smith and Price, 1973, Maynard Smith, 1982] but we designate evolutionary aspects of equilibrium computation as out of scope of the present paper simply to facilitate a more concise exposition of our results. Having laid out our overarching strategy and approach, we should point out that our main techniques consist in two parts, namely, in *definitions* and *results* on the properties of the concepts they define.

1.3.2 Our main definitions

An important idea in this paper is captured in the notion of an *equalizer*, an equilibrium notion generalizing interior equilibria. This idea emerged as follows: Chen et al. [2009] point out that a simple (inefficient) algorithm to find an equilibrium of a (general) bimatrix game is by support enumeration and the solution of a corresponding linear (feasibility) program iteratively for each support until an equilibrium is found. Of course the number of possible supports is exponential and, therefore, this algorithm is inefficient in all games but those with a small number of pure strategies. But there is an important detail in their formulation that can easily go unnoticed:

Given some support, the linear program Chen et al. [2009] formulate equalizes the pure strategies in the support and restricts the pure strategies outside the support to yield payoffs that are not superior to the payoffs pure strategies in the support yield. The important observation is that the solution of this linear program can yield an equilibrium lying *outside* the given support wherein the linear feasibility program is invoked (in that some pure strategies can receive zero probability). The idea of an *equalizer* emerged in the effort to understand when this phenomenon occurs.

One very important “technical result” in this paper (which can be used to prove $\mathbf{P} = \mathbf{PPAD}$ as a simple exercise) is what happens to a symmetric bimatrix game if there is no equalizer. One algorithmically interesting answer is that the absence of an equalizer implies the existence of a weakly dominated pure strategy, which can be detected and eliminated in polynomial time [Conitzer and Sandholm, 2005] without missing out on the overall equilibrium computation objective.

1.3.3 Our main analytical techniques

Our proof of this property relies on a pair of lemmas on the behavior of a multiplicative weights dynamic in the vector field of a symmetric bimatrix game which are results that emerged in our effort to prove conditional convergent behavior (conditioned on the behavior of the payoff vector field near an equilibrium). We designate the latter results out of scope to remain concise.

1.4 RELATED WORK

Our work and contributions in this paper were inspired by the success and beauty of *online learning theory*, that is, *machine learning theory* though not in the (most popular) statistical framework but instead in the *regret minimization framework* viewing nature as a nonstochastic opponent. We apply online learning principles to the problem of Nash equilibrium computation.

1.4.1 Multiplicative weights

One of the earliest contributions in online learning is the “weighted majority algorithm” of Littlestone and Warmuth [1994]. Freund and Schapire [1997] extend this algorithm to the problems of “*online allocation*” as well as that of “*boosting*.” Arora et al. [2012] treat the history and applications of such algorithms to learning, game theory, and optimization rather comprehensively.

Freund and Schapire [1999] prove convergence of a multiplicative weights algorithm to the *value* of a zero-sum game—the algorithm is employed by one of the two player positions and the opponent in the other position is assumed to be arbitrary. We note their result does not imply convergence to the minimax equilibrium, but it furnishes an alternative proof of von Neumann’s *minimax theorem*.

Minimax equilibria in zero-sum games (such as the equilibrium of the rock-paper-scissors game) may “repel” the multiplicative weights dynamic (but not always). This is easy to prove using modified versions of analytical techniques employed in this paper (but the proof is omitted for brevity). We note their proof techniques are based on the relative entropy as in this paper.

Similar negative results that multiplicative weights are not convergent to Nash equilibria are obtained by [Daskalakis et al., 2010] and [Balkan et al., 2012]. We note that, as mentioned earlier, we do not use multiplicative weights to compute Nash equilibria but rather as a mathematical tool to “probe” (and characterize) the incentive (payoff) structure of strategic games.

1.4.2 Equilibrium computation and approximation

PPAD (Polynomial Parity Arguments on Directed graphs) is a complexity class introduced by Papadimitriou [1994] to capture the difficulty of finding Brouwer fixed points. The archetypical problem in this class is to find an unbalanced vertex in a directed graph given some other unbalanced vertex in the same graph. Such a vertex can always be found but the conjecture was that the amount of time required to find one is exponential. Since **PPAD** was introduced, various problems have been shown to be **PPAD**-complete such as that of computing an N -player Nash equilibrium [Daskalakis et al., 2009] and that of computing a 2-player Nash equilibrium [Chen et al., 2009]. Adding support to the conjecture that **PPAD** captures hard problems, the Lemke-Howson algorithm, generally regarded as the most efficient algorithm for computing a 2-player Nash equilibrium, was shown to run in exponential time in the worst case [Savani and von Stengel, 2004].

There are two important threads of results on equilibrium computation and approximation, namely, the approximation of Nash equilibria in bimatrix games and the approximation of KKT (Karush, Kuhn, and Tucker) points in quadratic programs:

(1) With respect to the first thread of results, under the assumption that $\mathbf{PPAD} \neq \mathbf{P}$ the most significant open question concerned the existence of a polynomial time approximation scheme (PTAS) for equilibrium computation in bimatrix games (as it was well-known [Chen et al., 2009] that a full polynomial time approximation scheme (FPTAS) would imply $\mathbf{P} = \mathbf{PPAD}$, which was considered unlikely). There are two notions of approximation, namely, *additive* approximations and *multiplicative* ones. The first quasi-polynomial algorithm for additively approximate Nash equilibria in bimatrix games is due to Lipton et al. [2003] based on the principle of looking for equilibria of small support. There is a rich literature on the additively and multiplicatively approximations of Nash equilibria. An interesting alternative approach, however, is to find an exact Nash equilibrium with high probability. Bárány et al. [2005] present an algorithm for finding a Nash equilibrium in random 2-player games that runs in polynomial time with high probability. This algorithm is based on searching for equilibria of small support. We note that the general methods employed in this thread of research depart significantly from our algorithmic and proof techniques:

The algorithm we present in this paper instead looks for equilibria of (possibly) *large* support (as equalizers are typically, but not always so, interior equilibria) and iteratively reduces the size of the game (by iterated elimination of weakly dominated pure strategies), checking for an equalizer in each step, until a single pure strategy remains (implying the game is dominance solvable). Alternatively, one may iteratively reduce the size of the basis game by iterated elimination until no further reduction is possible in which case our results imply the existence of an equalizer. As noted earlier, in contrast to other previous equilibrium computation efforts, our algorithm focuses on symmetric bimatrix games trying to find an equilibrium that is symmetric. (Given a bimatrix game an “equivalent” symmetric bimatrix game can be obtained by the Gale, Kuhn, and Tucker symmetrization method [Jurg et al., 1992].) We have not verified if the absence of an equalizer in general bimatrix games implies the existence of a weakly dominated pure strategy.

(2) With respect to the second thread of results symmetric equilibria in *doubly symmetric bimatrix games* (that is, symmetric bimatrix games wherein the payoff matrix is symmetric) are known to coincide with KKT points of a quadratic optimization problem whose objective function is $X \cdot SX$, where S is the symmetric payoff matrix and X is a mixed strategy, and whose constraint set is a standard (probability) simplex corresponding to the space of mixed strategies. The computation of KKT points in (general) quadratic programs is known to admit an FPTAS [Ye, 1998]. Our results in this paper imply a polynomial-time algorithm for computing KKT points of quadratic programs over the probability simplex. We believe that our results can be extended to general quadratic programs (over general *polytopes*) but we have not verified the conjecture.

1.5 OVERVIEW OF THE REST OF THE PAPER

Section 2 discusses preliminaries on bimatrix games and Nash equilibria. Section 3 gives our equilibrium computation algorithm. We subsequently move on to prove its correctness starting with a preliminary technical result on the structure of the best response correspondence in Section 4. Section 5 is concerned with the definition and properties of “multiplicative weights dynamics” where some powerful technical lemmas important in the main result of this paper are stated and proven. Section 6 proves our main result concerning an elementary property of a symmetric bimatrix game without an equalizer, which is used to furnish a simple proof that $\mathbf{P} = \mathbf{PPAD}$ in the same section. Finally, Section 7 concludes the paper.

2 Preliminaries on bimatrix games

2.1 BIMATRIX GAMES

A 2-player (bimatrix) game in normal form is specified by a pair of $n \times m$ matrices A and B , the former corresponding to the *row player* and the latter to the *column player*. If $B = A^T$, where A^T is the transpose of A , the game is called *symmetric*. A *mixed strategy* for the row player is a probability vector $P \in \mathbb{R}^n$ and a mixed strategy for the column player is a probability vector $Q \in \mathbb{R}^m$. The *payoff* to the row player of P against Q is $P \cdot AQ$ and that to the column player is $P \cdot BQ$. Let us denote the space of probability vectors for the row player by \mathbb{P} and the corresponding space for the column player by \mathbb{Q} . A Nash equilibrium of a 2-player game (A, B) is a pair of mixed strategies P^* and Q^* such that all unilateral deviations from these strategies are not profitable, that is, for all $P \in \mathbb{P}$ and $Q \in \mathbb{Q}$, we simultaneously have that

$$\begin{aligned} P^* \cdot AQ^* &\geq P \cdot AQ^* \\ P^* \cdot BQ^* &\geq P^* \cdot BQ. \end{aligned}$$

We denote the set of Nash equilibria of (A, B) by $NE(A, B)$.

2.2 SYMMETRIC BIMATRIX GAMES

Let (C, C^T) be a symmetric bimatrix game. We call a Nash equilibrium (P^*, Q^*) of (C, C^T) symmetric if $P^* = Q^*$. Let (X^*, X^*) be a symmetric equilibrium. Then we call X^* a symmetric equilibrium strategy of (C, C^T) . A well-known result due to Nash is that every symmetric N -person game, and, therefore, also every symmetric bimatrix game, has a symmetric equilibrium.

Given (C, C^T) , we denote its set of pure strategies by $\mathcal{K}(C, C^T) = \{1, \dots, n\}$. Pure strategies are denoted either as i or as E_i , a probability vector whose mass is concentrated in position i . We denote the probability simplex in \mathbb{R}^n , where n is the number of pure strategies of (C, C^T) , by $\mathbb{X}(C, C^T)$. We often refer to $\mathbb{X}(C, C^T)$ either as the probability simplex of (C, C^T) or as the *evolution space* of (C, C^T) , and it corresponds to the space of mixed strategies. We denote the set of Nash equilibria of (C, C^T) by $NE(C, C^T)$ and the set of symmetric Nash equilibrium strategies thereof by $NE^+(C, C^T)$. Let $X \in \mathbb{X}(C, C^T)$ and define

$$\mathcal{C}(X) \doteq \{i \in \{1, \dots, n\} | X(i) > 0\}.$$

We call $\mathcal{C}(X)$ the *carrier* or *support* of strategy $X \in \mathbb{X}(C, C^T)$. We call $X \in \mathbb{X}(C, C^T)$ *interior* if $\mathcal{C}(X) = \{1, \dots, n\}$. We denote the relative interior of $\mathbb{X}(C, C^T)$ (that is, the set of all interior strategies of $\mathbb{X}(C, C^T)$) by $\overset{\circ}{\mathbb{X}}(C, C^T)$ and its boundary $\bar{\mathbb{X}}(C, C^T)$. The *carrier game* of a strategy $X \in \mathbb{X}(C, C^T)$ is the symmetric subgame of (C, C^T) whose pure strategies correspond to $\mathcal{C}(X)$.

3 Equilibrium computation algorithm

Let us first give the algorithm we propose for finding a symmetric equilibrium in a symmetric bimatrix game. We first need to declare some subroutines:

3.1 EQUALIZER EQUILIBRIA (OR EQUALIZERS)

The first subroutine we need checks if a symmetric bimatrix game has an “equalizer” and if so computes one. The declaration of this subroutine is:

The input is the payoff matrix C of a symmetric bimatrix game (C, C^T) and the output is either NULL if there is no “equalizer” of (C, C^T) or otherwise the output is a mixed strategy of (C, C^T) corresponding to an “equalizer” of (C, C^T) . But let us precisely define what the subroutine checks for (or computes)—to that end we start with the definition of an “equalizer:”

Definition 1. $X^* \in \mathbb{X}(C, C^T)$ is called an equalizer if

$$\forall X \in \mathbb{X}(C, C^T) : (X^* - X) \cdot CX^* = 0.$$

We denote the set of equalizers of (C, C^T) by $\mathbb{E}(C, C^T)$.

Note that $\mathbb{E}(C, C^T) \subseteq NE^+(C, C^T)$. Equalizers generalize interior symmetric equilibrium strategies, as every such strategy is an equalizer, but there exist symmetric bimatrix games with a non-interior equalizer (for example, if a column of C is constant, the corresponding pure strategy of (C, C^T) is an equalizer of (C, C^T)). Note that an equalizer can be computed in polynomial time by solving the linear (feasibility) program (LP)

$$(CX)_1 = \dots = (CX)_n, \quad \sum_{i=1}^n X(i) = 1, \quad X \geq 0,$$

which we may equivalently write as

$$CX = c\mathbf{1}, \quad \mathbf{1}^T X = 1, \quad X \geq 0,$$

where $\mathbf{1}$ is a column vector of ones of appropriate dimension. We may write this problem as a standard LP as follows: Let

$$A \doteq \begin{bmatrix} C & -\mathbf{1} \\ \mathbf{1}^T & 0 \end{bmatrix} \text{ and } Y \doteq \begin{bmatrix} X \\ c \end{bmatrix}.$$

then

$$AY = \begin{bmatrix} C & -\mathbf{1} \\ \mathbf{1}^T & 0 \end{bmatrix} \begin{bmatrix} X \\ c \end{bmatrix} = \begin{bmatrix} CX - c\mathbf{1} \\ \mathbf{1}^T X \end{bmatrix},$$

and the standard form of our LP, assuming $C > 0$, is

$$\begin{bmatrix} CX - c\mathbf{1} \\ \mathbf{1}^T X \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix}, X \geq 0, c \geq 0$$

where $\mathbf{0}$ is a column vector of zeros of appropriate dimension.

3.2 STRICTLY AND WEAKLY DOMINATED STRATEGIES

The second subroutine we need checks if a pure strategy, say E_i , of a symmetric bimatrix game is “weakly dominated.” The declaration of this subroutine is:

This subroutine receives as a input a pure strategy $i \in \mathcal{K}(C, C^T)$ and returns as output either TRUE if the pure strategy is “weakly dominated” by some mixed strategy and FALSE otherwise. But let us precisely define the (well-known in game theory) notion of dominance and see how this subroutine can run in polynomial time (with the complexity of linear programming).

There are two versions of dominance the stronger of which is commonly referred to as “strict dominance.” Let us give the precise definition:

Definition 2. *We say that the pure strategy E_i of (C, C^T) is strictly dominated if*

$$\exists X \in \mathbb{X}(C, C^T) \forall Y \in \mathbb{X}(C, C^T) : E_i \cdot CY < X \cdot CY.$$

It is well known that eliminating a strictly dominated strategy cannot eliminate equilibria. Our primary concern though is with a weaker dominance notion which doesn’t have this property:

Definition 3. *We say that the pure strategy E_i of (C, C^T) is weakly dominated if*

$$\exists X \in \mathbb{X}(C, C^T) \forall Y \in \mathbb{X}(C, C^T) : E_i \cdot CY \leq X \cdot CY$$

and

$$\exists Y \in \mathbb{X}(C, C^T) : E_i \cdot CY < X \cdot CY.$$

If E_i is not weakly dominated we say it is undominated.

Unlike strict elimination, eliminating weakly dominated strategies can eliminate equilibria, however, every equilibrium of the surviving subgame is also an equilibrium of the basis game. (Note also that, in weak iterated elimination, the order of elimination matters in that the surviving equilibria depend on the order by which strategies are eliminated. This matter is rather tangential to the main subject and result of this paper although a careful treatment of the subject is certainly important. Note that checking if a pure strategy is weakly dominated (by some mixed strategy) admits a polynomial-time algorithm (with the complexity of linear programming) [Conitzer and Sandholm, 2005] (see also [Myerson, 1991] and for iterated dominance see also [Knuth et al., 1988]).

3.3 THE EQUILIBRIUM COMPUTATION ALGORITHM

We may now give the pseudocode for the equilibrium computation algorithm (as Algorithm 1) computing a symmetric equilibrium (that is, a Nash equilibrium wherein both players use the same strategy) in a symmetric bimatrix game—such an equilibrium is guaranteed to exist [Nash, 1951].

Algorithm 1 SYMMETRIC-2-PLAYER-NASH(C)

```
1: if SIZE( $C$ ) == 1 then
2:   return 1
3: else
4:   if ( $X^* = \text{EQUALIZER}(C)$ ) != NULL then
5:     return  $X^*$ 
6:   else
7:      $i = 1$ 
8:     while WEAKLY-DOMINATED( $i, C$ ) == FALSE do
9:        $i++$ 
10:    end while
11:     $Y^* = \text{SYMMETRIC-2-PLAYER-NASH}(C/\{E_i\})$ 
12:    return ZERO-PAD( $Y^*$ )
13:  end if
14: end if
```

Notice the structure of the code is recursive: The algorithm receives as input the payoff matrix C of a symmetric bimatrix game (C, C^T) and returns as output a symmetric equilibrium strategy of (C, C^T) . Let us briefly go through the pseudocode: If there is only one pure strategy (Line 1), 1 is returned as the solution, that is, a probability vector of one entry (Line 2). If the number of pure strategies is greater than one, the algorithm first checks (in Line 4) if (C, C^T) has an equalizer (an operation that can be performed by solving an LP as noted earlier). If there is an equalizer, the routine terminates with the equalizer as the output (Line 5). Otherwise the algorithm iteratively looks for a weakly dominated pure strategy (Lines 7 to 10) and recursive calls itself on the symmetric subgame that is obtained by removing the pure strategy from the set of pure strategies (Line 11). Padding with zeros (Line 12) ensures consistency of the recursive calls. Note there are other related variants of this specific (pseudocode) implementation that also work.

We may now proceed to prove the previous algorithm's correctness. We start off by giving some elementary and well-known definitions and propositions on "best responses" for completeness:

4 Best responses

In this section, we give basic background on best responses. Propositions 1, 2, and 3 are standard.

Definition 4. Let (A, B) be a bimatrix game and $\mathbb{P} \times \mathbb{Q}$ be the space of strategy profiles of (A, B) . Let $Q \in \mathbb{Q}$ be an arbitrary column-player strategy. We say $P^*(Q) \in \mathbb{P}$ is a best response to Q if

$$P^*(Q) \in \arg \max \{P \cdot AQ \mid P \in \mathbb{P}\}.$$

Best responses of the column player against row-player strategies are defined similarly.

We have the following characterization of Nash equilibria in bimatrix games.

Proposition 1. $(X^*, Y^*) \in \mathbb{P} \times \mathbb{Q}$ is a Nash equilibrium of (A, B) if and only if $X^* \in \mathbb{P}$ is a best response to $Y^* \in \mathbb{Q}$ and, similarly, Y^* is a best response to X^* .

Proof. Straightforward from the definition of a Nash equilibrium. □

An important elementary property of best responses is given by the (aforementioned) *best response condition* [von Stengel, 2007]:

Proposition 2. *Given (A, B) let A and B be $m \times n$ matrices, $\mathbb{P} \times \mathbb{Q}$ be the space of mixed strategy profiles of (A, B) , $X \in \mathbb{P}$, and $Y \in \mathbb{Q}$. Then X is a best response to Y if and only if*

$$\forall i = \{1, \dots, m\} : \left[X(i) > 0 \Rightarrow (AY)_i = (AY)_{\max} \equiv \max_{i=1}^m (AY)_i \right].$$

The best response condition implies the following proposition.

Proposition 3. *$(X^*, Y^*) \in NE(A, B)$ if and only if we simultaneously have:*

- (i) *for all pure strategies $E_i \in \mathbb{P}$ such that $i \in \mathcal{C}(X^*)$, E_i is a best response to Y^* and*
- (ii) *for all pure strategies $E_j \in \mathbb{Q}$ such that $j \in \mathcal{C}(Y^*)$, E_j is a best response to X^* ,*

where $\mathbb{P} \times \mathbb{Q}$ is the space of mixed strategies of (A, B) .

5 Multiplicative weights: Properties and lemmas

The next step in the algorithm’s correctness proof is a set of lemmas on “multiplicative weights dynamics.” The specific dynamic we use to “probe” the mathematical structure and properties of symmetric bimatrix games is based on Hedge [Freund and Schapire, 1997, 1999]. In this setting, Hedge induces the following dynamic:

$$T_i(X) = X(i) \cdot \frac{\exp\{\alpha E_i \cdot CX\}}{\sum_{j=1}^n X(j) \exp\{\alpha E_j \cdot CX\}}, \quad i = 1, \dots, n. \quad (1)$$

In this section, we give some properties of (1). We first show that $\mathbb{X}(C, C^T)$ is invariant under T in that T never escapes the space of mixed strategies of the game (C, C^T) in which it is invoked. Actually we also show that if T is started in the relative interior of $\mathbb{X}(C, C^T)$, T cannot escape $\mathring{\mathbb{X}}(C, C^T)$ in a finite number of steps. Then we characterize fixed points and prove some powerful lemmas. We also prove the promised “better response property” as Lemma 8 below.

5.1 INVARIANCE

Note that, in general, α (typically referred to as the *learning rate*) may depend on X —we make this explicit in some lemmas for clarity.

Lemma 1. *We have*

$$\forall \alpha(X) \geq 0 : X \in \mathbb{X}(C, C^T) \Rightarrow T(X) \in \mathbb{X}(C, C^T)$$

and

$$\forall \alpha(X) \geq 0 : X \in \mathring{\mathbb{X}}(C, C^T) \Rightarrow T(X) \in \mathring{\mathbb{X}}(C, C^T),$$

where T is as in (1).

Proof. We have that, given any $X \in \mathbb{X}(C, C^T)$,

$$\forall \alpha(X) \geq 0 : \sum_{i=1}^n T_i(X) = \left(\sum_{j=1}^n X(j) \exp \{ \alpha(X) E_j \cdot CX \} \right)^{-1} \cdot \sum_{i=1}^n X(i) \exp \{ \alpha(X) E_i \cdot CX \} = 1,$$

implying *positive invariance* of $\mathbb{X}(C, C^T)$. Furthermore,

$$\forall \alpha(X) \geq 0 : X(i) > 0 \Rightarrow T_i(X) > 0,$$

further implying *positive invariance* of the relative interior $\mathring{\mathbb{X}}(C, C^T)$. \square

5.2 FIXED POINTS OF HEDGE

One of game theory's principles is that “equilibrium” is synonymous to *Nash equilibrium*, defined as a fixed point of the best response correspondence. The first order principle of the multiplicative weights approach to game theory is, not surprisingly, a fixed point of the multiplicative weights dynamic. In this section, we show how the two concepts are related. We refer to Nash fixed points as “equilibria” and to multiplicative weights fixed points as “fixed points” simply for reasons of remaining backward compatible with the game-theoretic nomenclature.

Definition 5. We say $X^* \in \mathbb{X}(C, C^T)$ is a fixed point of (C, C^T) if it is a fixed point of (1). That is, if $X^* = T(X^*)$. We denote the set of fixed points of (C, C^T) by $FX(C, C^T)$. Note that

$$\forall i \in \mathcal{K}(C, C^T) : E_i \in FX(C, C^T).$$

We denote the non-pure fixed points of (C, C^T) by $FX^+(C, C^T)$.

We have the following characterization of fixed points.

Theorem 1. $X \in FX(C, C^T)$ if and only if, X is a pure strategy or otherwise

$$\forall i, j \in \mathcal{C}(X) : (CX)_i = (CX)_j.$$

Proof. First we show sufficiency, that is, if for all $i, j \in \mathcal{C}(X)$, $(CX)_i = (CX)_j$, then $T(X) = X$: Some of the coordinates of X are zero and some are positive. Clearly, the zero coordinates will not become positive after applying the map. Now, notice that, for all $i \in \mathcal{C}(X)$, $\exp \{ \alpha(X)(CX)_i \} = \sum_{j=1}^n X(j) \exp \{ \alpha(X)(CX)_j \}$. Therefore, $T(X) = X$.

Now we show necessity, that is, if X is a fixed point of Hedge, then for all i and for all $j \in \mathcal{C}(X)$, $(CX)_i = (CX)_j$: Let $\hat{X}(i) = T_i(x)$. Because X is a fixed point, $\hat{X}(i) = X(i)$. Therefore,

$$\begin{aligned} \hat{X}(i) &= X(i) \\ \frac{X(i) \exp \{ \alpha(X)(CX)_i \}}{\sum_j X(j) \exp \{ \alpha(X)(CX)_j \}} &= X(i) \\ \exp \{ \alpha(X)(CX)_i \} &= \sum_j X(j) \exp \{ \alpha(X)(CX)_j \}, \end{aligned}$$

which implies

$$\exp \{ \alpha(X)((CX)_i - (CX)_j) \} = 1, X(i) > 0,$$

and, thus,

$$(CX)_i = (CX)_j, X(i) > 0.$$

This completes the proof. \square

The previous theorem says that $X^* \in FX(C, C^T)$ if and only if X^* is an equalizer of its carrier game. This is a necessary condition for equilibrium but it is not sufficient:

Lemma 2. *If $X^* \in \overset{\circ}{\mathbb{X}}(C, C^T)$, $X^* \in FX(C, C^T)$ if and only if $X^* \in NE^+(C, C^T)$.*

Proof. The lemma is a straightforward implication of the aforementioned *best response condition* and Theorem 1: Suppose the interior $X^* \in NE^+(C, C^T)$ and recall the best response condition: X is a best response to Y if and only if $X(i) > 0 \Rightarrow (CY)_i = (CY)_{\max}$. But a symmetric Nash equilibrium strategy is a best response to itself, therefore, $X^*(i) > 0 \Rightarrow (CX^*)_i = (CX^*)_{\max}$. Since X^* is interior, for all $i \in \mathcal{K}(C, C^T)$, $X^*(i) > 0$, therefore,

$$\forall i, j \in \mathcal{C}(X^*) : (CX^*)_i = (CX^*)_j.$$

The characterization of fixed points in Theorem 1 completes the proof. \square

As a straightforward implication of the previous lemma we obtain:

Lemma 3. *If $X^* \in FX(C, C^T)/NE^+(C, C^T)$, then $X^* \in \bar{\mathbb{X}}(C, C^T)$. That is, every non-equilibrium fixed point is a boundary strategy.*

Proof. If $X^* \notin \bar{\mathbb{X}}(C, C^T)$, then $X^* \in \overset{\circ}{\mathbb{X}}(C, C^T)$. \square

Overall we obtain:

Corollary 1. $NE^+(C, C^T) \subseteq FX(C, C^T)$.

Proof. Every symmetric equilibrium strategy is an equalizer of the its corresponding carrier game by Lemma 2. \square

The fundamental difference between fixed points and equilibria is that fixed points survive “inversion” of incentives. We call this principle “the virtue of vice versa” and it is proven next:

Theorem 2. $FX(C, C^T) = FX(-C, -C^T)$.

Proof. The theorem clearly holds for all pure fixed points. Furthermore, by Theorem 1, $X^* \in FX^+(C, C^T)$ if and only if

$$\forall X \in \mathbb{X}(C^*, C^{*T}) : (X^* - X) \cdot C^* X^* = 0$$

where C^* is X^* ’s carrier game. This completes the proof. \square

We do not invoke the virtue of vice versa in this paper, but it is important to keep in mind.

5.3 SOME LEMMAS

5.3.1 Preliminaries: Relative entropy (or Kullback-Leibler divergence)

Our analysis of multiplicative weights dynamics relies on the relative entropy function between probability distributions (also called *Kullback-Leibler divergence*). The relative entropy between the $n \times 1$ probability vectors $P > 0$ (that is, for all $i = 1, \dots, n$, $P(i) > 0$) and $Q > 0$ is given by

$$RE(P, Q) \doteq \sum_{i=1}^n P(i) \ln \frac{P(i)}{Q(i)}.$$

However, this definition can be relaxed: The relative entropy between $n \times 1$ probability vectors P and Q such that, given P , for all $Q \in \{Q \in \mathbb{X} | \mathcal{C}(P) \subset \mathcal{C}(Q)\}$, is

$$RE(P, Q) \doteq \sum_{i \in \mathcal{C}(P)} P(i) \ln \frac{P(i)}{Q(i)}.$$

We note the standard well-known properties of the relative entropy [Weibull, 1995, p.96] that (i) $RE(P, Q) \geq 0$, (ii) $RE(P, Q) \geq \|P - Q\|^2$, where $\|\cdot\|$ is the Euclidean distance, (iii) $RE(P, P) = 0$, and (iv) $RE(P, Q) = 0$ iff $P = Q$. Note (i) follows from (ii) and (iv) follows from (ii) and (iii).

5.3.2 The convexity lemma

Lemma 4. *Let T be as in (1). Then*

$$\forall X \in \mathring{\mathbb{X}}(C, C^T) \forall Y \in \mathbb{X}(C, C^T) : RE(Y, T(X)) \text{ is a convex function of } \alpha \geq 0.$$

Furthermore, unless $X \in FX(C, C^T)$, $RE(Y, T(X))$ is a strictly convex function of $\alpha > 0$.

Proof. We have

$$\begin{aligned} \frac{d}{d\alpha} RE(Y, \hat{X}) &= \\ &= \frac{d}{d\alpha} \left(\sum_{i \in \mathcal{C}(Y)} Y(i) \ln \left(\frac{Y(i)}{\hat{X}(i)} \right) \right) \\ &= \frac{d}{d\alpha} \left(\sum_{i \in \mathcal{C}(Y)} Y(i) \ln \left(Y(i) \cdot \frac{\sum_{j=1}^n X(j) \exp\{\alpha(CX)_j\}}{X(i) \exp\{\alpha(CX)_i\}} \right) \right) \\ &= \frac{d}{d\alpha} \left(\sum_{i \in \mathcal{C}(Y)} Y(i) \ln \left(\frac{\sum_{j=1}^n X(j) \exp\{\alpha(CX)_j\}}{X(i) \exp\{\alpha(CX)_i\}} \right) \right) \\ &= \sum_{i \in \mathcal{C}(Y)} Y(i) \frac{d}{d\alpha} \left(\ln \left(\frac{\sum_{j=1}^n X(j) \exp\{\alpha(CX)_j\}}{X(i) \exp\{\alpha(CX)_i\}} \right) \right). \end{aligned}$$

Furthermore, using $(\cdot)'$ as alternative notation (abbreviation) for $d/d\alpha(\cdot)$,

$$\frac{d}{d\alpha} \left(\ln \left(\frac{\sum_{j=1}^n X(j) \exp\{\alpha(CX)_j\}}{X(i) \exp\{\alpha(CX)_i\}} \right) \right) = \frac{X(i) \exp\{\alpha(CX)_i\}}{\sum_{j=1}^n X(j) \exp\{\alpha(CX)_j\}} \left(\frac{\sum_{j=1}^n X(j) \exp\{\alpha(CX)_j\}}{X(i) \exp\{\alpha(CX)_i\}} \right)'$$

and

$$\begin{aligned} \left(\frac{\sum_{j=1}^n X(j) \exp\{\alpha(CX)_j\}}{X(i) \exp\{\alpha(CX)_i\}} \right)' &= \frac{\sum_{j=1}^n X(j)(CX)_j \exp\{\alpha(CX)_j\} X(i) \exp\{\alpha(CX)_i\}}{(X(i) \exp\{\alpha(CX)_i\})^2} - \\ &\quad - \frac{X(i)(CX)_i \sum_{j=1}^n X(j) \exp\{\alpha(CX)_j\} \exp\{\alpha(CX)_i\}}{(X(i) \exp\{\alpha(CX)_i\})^2} = \\ &= \frac{\sum_{j=1}^n X(j)(CX)_j \exp\{\alpha(CX)_j\}}{X(i) \exp\{\alpha(CX)_i\}} - \\ &\quad - \frac{(CX)_i \sum_{j=1}^n X(j) \exp\{\alpha(CX)_j\}}{X(i) \exp\{\alpha(CX)_i\}}. \end{aligned}$$

Therefore,

$$\frac{d}{d\alpha} RE(Y, \hat{X}) = \frac{\sum_{j=1}^n X(j)(CX)_j \exp\{\alpha(CX)_j\}}{\sum_{j=1}^n X(j) \exp\{\alpha(CX)_j\}} - Y \cdot CX. \quad (2)$$

Furthermore,

$$\begin{aligned} \frac{d^2}{d\alpha^2} RE(Y, \hat{X}) &= \frac{\left(\sum_{j=1}^n X(j)((CX)_j)^2 \exp\{\alpha(CX)_j\}\right) \left(\sum_{j=1}^n X(j) \exp\{\alpha(CX)_j\}\right)}{\left(\sum_{j=1}^n X(j) \exp\{\alpha(CX)_j\}\right)^2} - \\ &\quad - \frac{\left(\sum_{j=1}^n X(j)((CX)_j) \exp\{\alpha(CX)_j\}\right)^2}{\left(\sum_{j=1}^n X(j) \exp\{\alpha(CX)_j\}\right)^2}. \end{aligned}$$

Jensen's inequality implies that

$$\frac{\sum_{j=1}^n X(j)((CX)_j)^2 \exp\{\alpha(CX)_j\}}{\sum_{j=1}^n X(j) \exp\{\alpha(CX)_j\}} \geq \left(\frac{\sum_{j=1}^n X(j)((CX)_j) \exp\{\alpha(CX)_j\}}{\sum_{j=1}^n X(j) \exp\{\alpha(CX)_j\}} \right)^2,$$

which is equivalent to the numerator of the second derivative being nonnegative as X is a probability vector. Note that the inequality is strict unless

$$\forall i, j \in \mathcal{C}(X) : (CX)_i = (CX)_j.$$

This completes the proof. \square

5.3.3 The stability lemmas

In the next lemma, we invoke the convexity lemma. We have:

Lemma 5. *Let $Y, X \in \mathbb{X}(C, C^T)$ such that $\mathcal{C}(Y) \subset \mathcal{C}(X)$ and such that $(Y - X) \cdot CX > 0$. Then there exists $\bar{\alpha} > 0$ (which may depend on Y and X) such that*

$$\forall \alpha \in (0, \bar{\alpha}) : RE(Y, \hat{X}) < RE(Y, X),$$

where $\hat{X} = T(X)$ and T is as in (1).

Proof. (2) implies that

$$\left. \frac{d}{d\alpha} RE(Y, \hat{X}) \right|_{\alpha=0} = (X - Y) \cdot CX < 0.$$

By the convexity $RE(Y, \hat{X})$ as a function of α either

$$\forall \alpha > 0 : RE(Y, \hat{X}) < RE(Y, X)$$

or

$$\exists \bar{\alpha} > 0 : RE(Y, \hat{X}) = RE(Y, X)$$

implying

$$\forall \alpha \in (0, \bar{\alpha}) : RE(Y, \hat{X}) < RE(Y, X)$$

as claimed. \square

Under an additional assumption we can show an even stronger (important in our main result) property:

Lemma 6. *Let X be an interior strategy in the probability simplex of (C, C^T) , assume $X \notin NE^+(C, C^T)$ and let*

$$Y \in \arg \max \{Y \cdot CX \mid Y \in \mathbb{X}(C, C^T)\}.$$

That is, Y is a best response to X . Then

$$\forall \alpha > 0 : RE(Y, \hat{X}) < RE(Y, X),$$

where $\hat{X} = T(X)$ and T is as in (1).

Proof. If $X \in NE^+(C, C^T)$, $\forall \alpha \geq 0 : RE(Y, \hat{X}) = RE(Y, X)$ since $\forall \alpha \geq 0 : \hat{X} = X$. Otherwise:

$$RE(Y, \hat{X}) - RE(Y, X) = \sum_{i \in \mathcal{C}(Y)} Y(i) \ln \left(\frac{\sum_{j=1}^n X(j) \exp \{\alpha (CX)_j\}}{\exp \{\alpha (CX)_i\}} \right) \quad (3)$$

$$= \ln \left(\sum_{j=1}^n X(j) \exp \{\alpha (CX)_j\} \right) - \alpha Y \cdot CX \quad (4)$$

$$= \ln \left(\sum_{j=1}^n X(j) \exp \{\alpha ((CX)_j - Y \cdot CX)\} \right) \quad (5)$$

$$\leq \sum_{j=1}^n X(j) \exp \{\alpha ((CX)_j - Y \cdot CX)\} - 1. \quad (6)$$

(3) follows by plugging in the definition of relative entropy and straight algebra. (4) follows by straightforward properties of the natural logarithm. (5) follows similarly. (6) follows from the standard property of the natural logarithm that $\ln(y) \leq y - 1, y > 0$.

If

$$\forall j \in \{1, \dots, n\} : (CX)_j = Y \cdot CX,$$

then $X \in NE^+(C, C^T)$ and, as noted earlier, in this case

$$\forall \alpha \geq 0 : RE(Y, \hat{X}) = RE(Y, X).$$

Assume, therefore, that

$$Y \cdot CX = \max_{j=1}^n (CX)_j > \min_{j=1}^n (CX)_j,$$

where the inequality on the left follows by the best response condition (Proposition 2). Then

$$\forall \alpha \geq 0 : \exp \{\alpha ((CX)_j - Y \cdot CX)\} \leq 1$$

and there exists $j \in \{1, \dots, n\}$ such that

$$\forall \alpha > 0 : \exp \{\alpha ((CX)_j - Y \cdot CX)\} < 1.$$

Therefore,

$$\forall \alpha > 0 : \sum_{j=1}^n X(j) \exp \{\alpha ((CX)_j - Y \cdot CX)\} < \sum_{j=1}^n X(j) = 1.$$

This completes the proof. □

5.3.4 An instability lemma

The following lemma is crucial. We can prove it in two ways, one invoking the aforementioned convexity lemma and the other by simply invoking Jensen's inequality. We show both proofs:

Lemma 7. *Let $Y, X \in \mathbb{X}(C, C^T)$ such that $\mathcal{C}(Y) \subseteq \mathcal{C}(X)$ and such that $Y \neq X$. If*

$$\max\{(CX)_i | i \in \mathcal{C}(X)\} > \min\{(CX)_i | i \in \mathcal{C}(X)\},$$

then

$$X \cdot CX - Y \cdot CX \geq 0 \Rightarrow \forall \alpha > 0 : RE(Y, \hat{X}) - RE(Y, X) > 0.$$

where $\hat{X} = T(X)$ and T is as in (1).

Proof. We have, by Jensen's inequality, that

$$\begin{aligned} RE(Y, \hat{X}) - RE(Y, X) &= \sum_{i \in \mathcal{C}(Y)} Y(i) \ln \frac{Y(i)}{\hat{X}(i)} - \sum_{i \in \mathcal{C}(Y)} Y(i) \ln \frac{Y(i)}{X(i)} \\ &= - \sum_{i \in \mathcal{C}(Y)} Y(i) \ln \hat{X}(i) + \sum_{i \in \mathcal{C}(Y)} Y(i) \ln X(i) \\ &= \sum_{i \in \mathcal{C}(Y)} Y(i) \ln \frac{X(i)}{\hat{X}(i)} \\ &= \sum_{i \in \mathcal{C}(Y)} Y(i) \ln \left(\frac{\sum_{j=1}^n X(j) \exp\{\alpha E_j \cdot CX\}}{\exp\{\alpha E_i \cdot CX\}} \right) \\ &= \ln \left(\sum_{j=1}^n X(j) \exp\{\alpha E_j \cdot CX\} \right) - \sum_{i \in \mathcal{C}(Y)} Y(i) \ln (\exp\{\alpha E_i \cdot CX\}) \\ &= \ln \left(\sum_{j=1}^n X(j) \exp\{\alpha E_j \cdot CX\} \right) - \alpha \sum_{i \in \mathcal{C}(Y)} Y(i) (E_i \cdot CX) \\ &\geq \alpha \sum_{j=1}^n X(j) (E_j \cdot CX) - \alpha \sum_{i \in \mathcal{C}(Y)} Y(i) (E_i \cdot CX) \\ &= \alpha (X \cdot CX - Y \cdot CX). \end{aligned}$$

Therefore,

$$X \cdot CX - Y \cdot CX \geq 0 \Rightarrow \forall \alpha > 0 : RE(Y, \hat{X}) - RE(Y, X) \geq 0.$$

If $\max_{i \in \mathcal{C}(X)} \{(CX)_i\} > \min_{i \in \mathcal{C}(X)} \{(CX)_i\}$, since Jensen's inequality is strict,

$$X \cdot CX - Y \cdot CX \geq 0 \Rightarrow \forall \alpha > 0 : RE(Y, \hat{X}) - RE(Y, X) > 0.$$

This completes the proof. □

Second proof of Lemma 7. By the convexity lemma (Lemma 4), unless $X \in FX(C, C^T)$, $RE(Y, \hat{X})$ is strictly convex. (2) implies that

$$\left. \frac{d}{d\alpha} RE(Y, \hat{X}) \right|_{\alpha=0} = (X - Y) \cdot CX \geq 0.$$

Noting that $RE(Y, \hat{X})|_{\alpha=0} = RE(Y, X)$ completes the proof. □

5.4 BETTER RESPONSE PROPERTY

Let us finally show our promised “better response property.”

Definition 6. We say that the map $T : \mathbb{X}(C, C^T) \rightarrow \mathbb{X}(C, C^T)$ is a better response dynamic if

$$\forall X \in \mathbb{X} : (\hat{X} - X) \cdot CX > 0,$$

where $\hat{X} = T(X)$, unless X is a fixed point of T .

The next lemma shows that T as defined in (1) is a better response dynamic.

Lemma 8. $X \in \mathring{\mathbb{X}}(C, C^T)$ implies $\hat{X} \cdot CX \equiv T(X) \cdot CX > X \cdot CX$ unless $X \in FX(C, C^T)$.

Proof. Suppose first that $X \notin NE^+(C, C^T)$, which implies $(CX)_{\max} > (CX)_{\min}$. Since the relative entropy between \hat{X} and X is positive:

$$RE(\hat{X}, X) = \sum_{i=1}^n \hat{X}(i) \ln \left(\frac{\hat{X}(i)}{X(i)} \right) = \sum_{i=1}^n \hat{X}(i) \ln \left(\frac{\exp \{ \alpha(X) E_i \cdot CX \}}{\sum_{j=1}^n X(j) \exp \{ \alpha(X) E_j \cdot CX \}} \right) > 0.$$

Therefore,

$$\ln \left(\sum_{j=1}^n X(j) \exp \{ \alpha(X) E_j \cdot CX \} \right) < \sum_{i=1}^n \hat{X}(i) \ln (\exp \{ \alpha(X) E_i \cdot CX \}) = \alpha(X) \hat{X} \cdot CX. \quad (7)$$

Furthermore, since $\ln(\cdot)$ is concave, we obtain by Jensen’s inequality that

$$\alpha X \cdot CX < \ln \left(\sum_{j=1}^n X(j) \exp \{ \alpha(X) E_j \cdot CX \} \right). \quad (8)$$

(7) and (8) imply that

$$(\hat{X} - X) \cdot CX > 0$$

as desired. If $X \in FX(C, C^T)$, Theorem 1 implies $\hat{X} = X$. \square

6 Correctness proof of the equilibrium computation algorithm

In this section, we prove the correctness of Algorithm 1 and $\mathbf{P} = \mathbf{PPAD}$. We need the following result on strategic dominance due to Pearce [1984] (see also [Weibull, 1995, p. 14]):

Proposition 4. The pure strategy E_i of (C, C^T) is undominated if and only if

$$\exists X \in \mathring{\mathbb{X}}(C, C^T) : E_i \text{ is a best response to } X.$$

We also need the following lemmas:

Lemma 9. Let $Y \in \bar{\mathbb{X}}(C, C^T)$ and $X \in \mathring{\mathbb{X}}(C, C^T)$. If $RE(Y, T(X))$ is strictly convex, then exactly one of the following hold:

$$(i) \lim_{\alpha \rightarrow \infty} (RE(Y, T(X)) - RE(Y, X)) = +\infty.$$

$$(ii) \lim_{\alpha \rightarrow \infty} (RE(Y, T(X)) - RE(Y, X)) < 0.$$

Furthermore, a necessary and sufficient condition for (ii) is that Y is a best response to X .

Proof. By the “convexity lemma” (Lemma 4), $RE(Y, T(X))$ is strictly convex if and only if X is not a fixed point. So let us assume this is the case. There are now three possibilities:

First possibility: Assume $(Y - X) \cdot CX \leq 0$. Then the “instability lemma” (Lemma 7) implies

$$\lim_{\alpha \rightarrow \infty} (RE(Y, T(X)) - RE(Y, X)) = +\infty.$$

Second possibility: Assume $(Y - X) \cdot CX > 0$ but Y is not a best response to X . Recall (5):

$$RE(Y, \hat{X}) - RE(Y, X) = \ln \left(\sum_{j=1}^n X(j) \exp \{ \alpha ((CX)_j - Y \cdot CX) \} \right)$$

and note that by the assumption that Y is not a best response to X there exists i such that $(CX)_i > Y \cdot CX$ implying that

$$\lim_{\alpha \rightarrow \infty} \left(\sum_{j=1}^n X(j) \exp \{ \alpha ((CX)_j - Y \cdot CX) \} \right) = +\infty.$$

Therefore,

$$\lim_{\alpha \rightarrow \infty} (RE(Y, T(X)) - RE(Y, X)) = +\infty.$$

Third possibility: Assume Y is a best response to X (implying $(Y - X) \cdot CX > 0$). Then Lemma 6 implies that

$$\lim_{\alpha \rightarrow \infty} (RE(Y, T(X)) - RE(Y, X)) < 0.$$

This completes the proof. □

Lemma 10. Suppose $\mathring{\mathbb{X}}(C, C^T) \cap FX(C, C^T) = \emptyset$, let $X \in \mathring{\mathbb{X}}$, and

$$Y \in \arg \max \{ \mathcal{Y} \cdot CX \mid \mathcal{Y} \in \mathbb{X}(C, C^T) \}.$$

That is, Y is a best response to X . Then

$$\forall X \in \mathring{\mathbb{X}}(C, C^T) : Y \in \arg \max \{ \mathcal{Y} \cdot CX \mid \mathcal{Y} \in \mathbb{X}(C, C^T) \}.$$

Proof. It is well-known that the best response correspondence is upper hemicontinuous. Finally, upper hemicontinuity of the best response correspondence implies that the set of strategies against which Y is a best response (that is, the lower inverse of Y under the best response correspondence) is closed [Corbae et al., 2009, Theorem 6.1.23]. Suppose now Y is a best response to an interior strategy and that the boundary of the set of strategies against which Y is a best response intersects the interior $\mathring{\mathbb{X}}(C, C^T)$ of the probability simplex of (C, C^T) . Let X lie on the boundary of the intersection. Then there exists $X' \in \mathring{\mathbb{X}}(C, C^T)$ such that

$$\forall \epsilon \in (0, 1] : Y \notin \arg \max \{ \mathcal{Y} \cdot CX(\epsilon) \mid \mathcal{Y} \in \mathbb{X}(C, C^T) \}$$

where

$$X(\epsilon) = (1 - \epsilon)X + \epsilon X'.$$

That is, Y is best response against X but it is not a best response against any other $X(\epsilon)$. Since, for all $X \in \overset{\circ}{\mathbb{X}}(C, C^T)$, $X \notin FX(C, C^T)$, for all $X \in \overset{\circ}{\mathbb{X}}(C, C^T)$, $RE(Y, T(X))$ is strictly convex (by the convexity lemma). Lemma 9 implies then that

$$\lim_{\alpha \rightarrow \infty} (RE(Y, T(X)) - RE(Y, X)) < 0$$

and

$$\forall \epsilon \in (0, 1] : \lim_{\alpha \rightarrow \infty} (RE(Y, T(X(\epsilon))) - RE(Y, X(\epsilon))) = +\infty.$$

Since $RE(Y, T(X(\epsilon))) - RE(Y, X(\epsilon))$ is continuous, this is clearly impossible. Hence the lemma. \square

Lemma 11. Suppose $\overset{\circ}{\mathbb{X}}(C, C^T) \cap FX(C, C^T) = \emptyset$, let $X \in \overset{\circ}{\mathbb{X}}$, and

$$Y \in \arg \max \{Y \cdot CX \mid Y \in \overset{\circ}{\mathbb{X}}(C, C^T)\}.$$

That is, Y is a best response to X . Then

$$\forall X \in \overset{\circ}{\mathbb{X}}(C, C^T) : (Y - X) \cdot CX > 0.$$

Proof. Since $X \notin NE^+(C, C^T)$, the best response condition (Proposition 2) implies that

$$Y \in \arg \max \{Y \cdot CX \mid Y \in \overset{\circ}{\mathbb{X}}(C, C^T)\} \Rightarrow (Y - X) \cdot CX > 0$$

and, then, Lemma 10 completes and proof. \square

Using the previous lemma we obtain the following theorem, which is a statement about the incentive structure of a symmetric bimatrix game without an (interior) equalizer:

Theorem 3. $\overset{\circ}{\mathbb{X}}(C, C^T) \cap FX(C, C^T) = \emptyset$ if and only if there exists a weakly dominated pure strategy of (C, C^T) .

Proof. By Proposition 4, the pure strategy E_i is undominated if and only if it is a best response to some interior (completely mixed) strategy. Suppose, for the sake of contradiction that every pure strategy is a best response to some interior strategy. Then Lemma 11 implies

$$\forall i \in \mathcal{K}(C, C^T) \forall X \in \overset{\circ}{\mathbb{X}}(C, C^T) : (E_i - X) \cdot CX > 0. \quad (9)$$

We consider two possibilities:

First possibility: Assume first $\mathcal{K}(C, C^T) \geq 3$: Consider a pure strategy E_i , a pair of pure strategies $E_{j_1} \neq E_{j_2}$ different from E_i , and let $Y = (1 - \epsilon)E_{j_1} + \epsilon E_{j_2}$, $\epsilon \in (0, 1)$, such that there exists a convex combination of E_i and Y in $\overset{\circ}{\mathbb{X}}(C, C^T)$. Then all such convex combinations are interior. Furthermore (9) implies

$$\forall X \text{ as convex combinations of } E_i \text{ and } Y : (E_i - X) \cdot CX > 0 \text{ and } (Y - X) \cdot CX > 0,$$

which is clearly an impossibility as $(E_i - X) \cdot CX > 0$ implies the angle between the payoff vector CX and $E_i - X$ is acute and $(Y - X) \cdot CX > 0$ similarly. But $E_i - X$ and $Y - X$ point in exactly opposite directions leading to a contradiction. Therefore, there exists a weakly dominated pure strategy.

Second possibility: If $\mathcal{K}(C, C^T) = 2$, the argument is entirely analogous.

The reverse direction is straightforward from Propositions 3 and 4: If $\mathring{\mathbb{X}}(C, C^T) \cap FX(C, C^T) \neq \emptyset$, every pure strategy is a best response to the interior symmetric equilibrium strategy by Proposition 3 (on best responses). Therefore, Proposition 4 implies every pure strategy is undominated. \square

Based on the previous results, we have:

6.1 CORRECTNESS PROOF OF ALGORITHM 1

To prove the correctness of Algorithm 1 it suffices to prove that if $\mathbb{E}(C, C^T) = \emptyset$ there exists a weakly dominated pure strategy. But this is a simple implication of Theorem 3 and the straightforward observation that

$$\mathbb{E}(C, C^T) = \emptyset \Rightarrow \mathring{\mathbb{X}}(C, C^T) \cap FX(C, C^T) = \emptyset.$$

Note the algorithm clearly runs in polynomial time.

6.2 $\mathbf{P} = \mathbf{PPAD}$

The problem of computing an equilibrium of a bimatrix game is **PPAD**-complete [Daskalakis et al., 2009, Chen et al., 2009]. Furthermore equilibrium computation in a bimatrix game can be reduced to the problem of equilibrium computation in a symmetric bimatrix game by the symmetrization method of [Jurg et al., 1992]. Since Algorithm 1 always terminates with a symmetric equilibrium of the symmetric game whose payoff matrix it receives as input, and since it is straightforward to verify that its running time is polynomial, that $\mathbf{P} = \mathbf{PPAD}$ follows immediately.

7 Concluding remarks

The computational complexity of our algorithm is good but it is certainly not “perfect” as it can be improved by a significant factor. Improving this algorithm and extending our results to other computational complexity classes is the main subject of our ongoing work on this topic.

Acknowledgments

I would like to thank Dick Lipton for comments on a previous draft. Part of this work was performed as I was an associate professor at the JST Erato Kawarabayashi project at the National Institute of Informatics in Tokyo, Japan. I colloquially call the multiplicative weights lemmas “Hatori Hanzo” lemmas as a tribute to the country that hosted their development.

References

S. Arora, E. Hazan, and S. Kale. The multiplicative weights update method: A meta-algorithm and its applications. *Theory of Computing*, 8:121–164, 2012.

- M. Balkan, F. Constantin, and R. Mehta. The weighted majority algorithm does not converge in nearly zero-sum games. In *Proc. ICML 2012 Workshop on Markets, Mechanisms, and Multi-Agent Models*, 2012.
- I. Bárány, S. Vempala, and A. Vetta. Nash equilibria in random games. In *Proc. 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 123–131, 2005.
- X. Chen, X. Deng, and S. Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3), 2009.
- V. Conitzer and T. Sandholm. Complexity of (iterated) dominance. In *Proc. of the 6th ACM conference on Electronic Commerce (EC '05)*, pages 88–97, 2005.
- D. Corbae, M. B. Stinchcombe, and J. Zeman. *An Introduction to Mathematical Analysis for Economic Theory and Econometrics*. Princeton University Press, Princeton, NJ, 2009.
- C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM J. Comput.*, 39(1):195–259, 2009.
- C. Daskalakis, R. Frongillo, C. H. Papadimitriou, G. Pierrakos, and G. Valiant. On learning algorithms for Nash equilibria. In *Proc. 3rd International Symposium on Algorithmic Game Theory (SAGT 2010)*, 2010.
- Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.
- Y. Freund and R. E. Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29:79–103, 1999.
- A. P. Jurg, M. J. M. Jansen, J. A. M. Potters, and S. H. Tijs. A symmetrization for finite two-person games. *ZOR – Methods and Models for Operations Research*, 36:111–123, 1992.
- D. E. Knuth, C. H. Papadimitriou, and J. N. Tsitsiklis. A note on strategy elimination in bimatrix games. *Operations Research Letters*, 7(3):103–107, 1988.
- R. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *Proc. EC'03*, pages 36–41, 2003.
- N. Littlestone and M. K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108:212–261, 1994.
- J. Maynard Smith. *Evolution and the Theory of Games*. Cambridge University Press, 1982.
- J. Maynard Smith and G. R. Price. The logic of animal conflict. *Nature*, 246:15–18, 1973.
- R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.
- J. F. Nash. Non-cooperative games. *The Annals of Mathematics, Second Series*, 54(2):286–295, Sept. 1951.
- C. H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Sciences*, 48(3):498–532, 1994.
- D. G. Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050, 1984.

- R. Savani and B. von Stengel. Exponentially many steps for finding a Nash equilibrium in a bimatrix game. In *Proc. 45th Annual IEEE Symposium on Foundations of Computer Science*, pages 258–267, 2004.
- B. von Stengel. Equilibrium computation for two-player games in strategic and extensive form. In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 3, pages 53–78. Cambridge University Press, 2007.
- J. W. Weibull. *Evolutionary Game Theory*. MIT Press, 1995.
- Y. Ye. On the complexity of approximating a KKT point of quadratic programming. *Mathematical Programming*, 80:195–211, 1998.